arXiv:1703.00197v1 [math.GR] 1 Mar 2017

Minimal and canonical images

Christopher Jefferson, Eliza Jonauskyte, Markus Pfeiffer

University of St Andrews
School of Computer Science
North Haugh
St Andrews
KY16 9SX
Scotland

Rebecca Waldecker

Martin-Luther-Universität Halle-Wittenberg
Institut für Mathematik
06099 Halle
Germany

Abstract

We describe a family of new algorithms for finding the canonical image of a set of points under the action of a permutation group. This family of algorithms makes use of the orbit structure of the group, and a chain of subgroups of the group, to efficiently reduce the amount of search which must be performed to find a canonical image.

We present both a formal proof of correctness of our algorithms and experiments on different permutation groups, which compare our algorithms with the previous state of the art.

Keywords: Minimal Images, Canonical Images, Computation, Group

Email addresses: caj21@st-andrews.ac.uk (Christopher Jefferson), ej31@st-andrews.ac.uk (Eliza Jonauskyte), markus.pfeiffer@st-andrews.ac.uk (Markus Pfeiffer), rebecca.waldecker@mathematik.uni-halle.de (Rebecca Waldecker)

URL: http://caj.host.cs.st-andrews.ac.uk/ (Christopher Jefferson), https://www.morphism.de/~markusp/ (Markus Pfeiffer), http://conway1.mathematik.uni-halle.de/~waldecker/index-english.html (Rebecca Waldecker)

Theory, Permutation Groups.

1. Background

Many combinatorial and group theory problems [14, 4, 2] are equivalent to finding, given a group G which acts on a finite set Ω and a subset $X \subseteq \Omega$, a partition of X into subsets which are in the same orbit of G.

We can solve such problems by taking two elements of X, and searching for an element of G which maps one to the other. However, this requires a possible $O(|X|^2)$ checks, if all elements of X are in different orbits.

Given a group G acting on a set Ω, a canonical function maps each element of Ω to a distinguished element of its orbit under G. Using a canonical function we can check if two members of Ω are in the same orbit by applying the canonical function to both and checking if the results are equal. More importantly, we can solve the problem of partitioning X into orbit-equivalent subsets by performing $O(|X|)$ canonical image calculations, and then comparing the resulting canonical images.

The canonical image problem has a long history. Jeffrey Leon [9] discusses three types of problems on permutation groups – subgroup-type problems (finding the intersection of several groups), coset-type problems (deciding whether or not the intersection of a series of cosets is empty, and if not, finding their intersection) and canonical-representative-type problems. He claims to have an algorithm to efficiently solve the canonical-representative problem, but does not discuss it further. His comments have inspired mathematicians and computer scientists to work on questions related to minimal images and canonical images.

One of the most well-studied canonical-image problems is the canonical graph problem. Current practical systems derive from the partition refinement techniques, which were first practically used for graph automorphisms by McKay [11] in the `Nauty` system. There have been a series of improvements to this technique, including `Saucy` [1], `Bliss` [8] and `Traces` [12].

We cannot, however, directly apply the existing work for graph isomorphism to finding canonical images in arbitrary groups. This is because McKay's Graph Isomorphism algorithm only considers finding the canonical image of

a graph under the action of the full symmetric group on the set of vertices. Many applications require finding canonical images under the action of subgroups of the full symmetric group.

One example of a canonicalizing function is, given a total ordering on X, to map each value of X to the smallest element in its orbit under G. This *Minimal image problem* has been treated by Linton in [10]. Pech and Reichard [13] apply techniques similar to Linton's to enumerate orbit representatives of subsets of Ω under the action of a permutation group on Ω. Linton gives a practical algorithm for finding the smallest image of a set under the action of a given permutation group. Our new algorithm, inspired by Linton's work, is designed to find canonical images: we extend and generalise Linton's technique using a new orbit-based counting technique. In this paper we first introduce some notation and explain the concepts that go into the algorithm, then we prove the necessary results and finish with experiments that demonstrate how this new algorithm is superior to the previously published techniques.

2. Minimal and Canonical Images

Throughout this paper, Ω will be a finite set, G a subgroup of $\mathrm{Sym}(\Omega)$, and Ω will be ordered by some (not necessarily total) order $\leqslant$.

If $\alpha \in \Omega$, then we denote the orbit of α under G by α^G.

Definition 2.1. *A **canonical function** C for the action of G on Ω is a function $C : \Omega \to \Omega$ such that, for all $\omega \in \Omega$, it is true that*

- $C(\omega) \in \omega^G$, *and*
- $C(\omega^g) = C(\omega)$ *for all* $g \in G$.

*In this situation we call $C(\omega)$ the **canonical image** of $\omega \in \Omega$ (with respect to G in this particular action).*

*For all $\omega \in \Omega$ we say that $g_\omega \in G$ is a **canonicalizing element** if and only if $\omega^{g_\omega} = C(\omega)$.*

One way to view canonical images is to pick a well-defined representative from each orbit of G on Ω.

In practice we want to be able to find canonical images effectively and efficiently. In some situations we are interested in computing a canonicalizing element, which might not be uniquely determined. Our algorithms will always produce a canonicalizing element as a byproduct of search, and it is trivial to extract. We choose to make this explicit here to make the exposition clearer.

Minimal images are a special type of canonical image.

Remark 2.2. *Suppose that $\preccurlyeq$ is a partial order on Ω such that any two elements in the same orbit can be compared by $\preccurlyeq$.*

Then the function $\mathrm{Min}_{\preccurlyeq}$, which for all $\omega \in \Omega$ maps ω to the smallest element in its orbit, is a canonical function.

In practical applications we are interested in more structure, namely in structures that G can act on naturally via the action on a given set Ω. These structures include subsets of Ω, graphs with vertex set Ω, sets of maps with domain or range Ω, and so on.

In this paper, our main application will be finding canonical images when acting on a set of subsets of Ω.

Definition 2.3. *Suppose that $\preccurlyeq$ is a total order of Ω. Then we introduce a total order $\preccurlyeq$ on $\mathcal{P}(\Omega)$ as follows:*

*We say that A **is less than** B if and only if A contains an element a such that $a \notin B$ and $a \preccurlyeq b$ for all $b \in B \backslash A$.*

Example 2.4. *Let $\Omega := \{1, 2, 3, 4, 5, 6, 7\}$ with the natural order and let $A := \{1, 3, 4\}$, $B := \{3, 5, 7\}$ and $C := \{3, 6, 7\}$.*

Now A is less than B, because $1 \in A$, $1 \notin B$ and 1 is smaller than all the elements in B, in particular those not in A. Moreover A is less than C for the same reason. Furthermore, B is less than C, because $5 \in B$, $5 \notin C$, and if we look at $C \backslash B$, then this only contains the element 6 and 5 is smaller.

If G is a subgroup of $\mathcal{S}_\Omega$, and $\omega \in \Omega$ then we denote by G_ω the point stabilizer of ω in G. For distinct elements $x, y \in \Omega$, we denote by $G_{x \mapsto y}$ the set of all elements of G that map x to y. This set may be empty.

We remark that the above information is readily available from a stabilizer chain for the group G, which can be calculated efficiently. For further details we refer the reader to [5]. We now introduce some notation and then prove a basic result about cosets.

Definition 2.5. *Let G be a permutation group acting on a totally ordered set $(\Omega, \leqslant)$, and let $\preccurlyeq$ denote the induced ordering as explained in Definition 2.3. Let H be a subgroup of G and $S \subseteq \Omega$. Then we define* ***the minimal image of*** *S* ***under*** *H to be the smallest element in the set $\{S^h | h \in H\}$ with respect to $\preccurlyeq$.*

In order to simplify notation, we will from now on write $\leqslant$ for the induced order and then we write $\mathrm{Min}(H, S, \leqslant)$ for the minimal image of S under H.

Lemma 2.6. *Let G be a permutation group acting on a totally ordered set $(\Omega, \leqslant)$, and let H be a subgroup of G and $S \subseteq \Omega$. Then the following hold for all $x, y \in \Omega$:*

(i) *For all $\sigma \in H_{x \mapsto y}$ it is true that $\sigma \cdot H_y = H_{x \mapsto y} = H_x \cdot \sigma$.*
(ii) *If $\sigma \in H_{x \mapsto y}$, then $\mathrm{Min}(\sigma \cdot H, S, \leqslant) = \mathrm{Min}(H, S^\sigma, \leqslant)$.*

Proof. If $\sigma \in H_{x \mapsto y}$, then multiplication by σ from the right or left is a bijection on H, respectively. For all $\alpha \in H_x$ we have that $\alpha \cdot \sigma$ maps x to y and for all $\beta \in H_y$ we see that $\sigma \cdot \beta$ also maps x to y. This implies the first statement. For (ii) we just look at the definition:

$\mathrm{Min}(\sigma \cdot H, S, \leqslant)$ denotes the smallest element in the set $\{S^{\sigma \cdot h} \mid h \in H\}$ and $\mathrm{Min}(H, S^\sigma, \leqslant)$ denotes the smallest element in the set $\{(S^\sigma)^h \mid h \in H\}$, which is the same set. □

2.1. Worked Example

We will find minimal, and later canonical, images using similar techniques to Linton in [10]. This algorithm splits the problem into small sub-problems, by splitting a group into the cosets of a point stabilizer. We will begin by demonstrating this general technique with a worked example.

Example 2.7. *In the following example we will look at $\Omega = \{1, 2, 3, 4, 5, 6\}$, the subgroup $G = \langle (14)(23)(56), (126) \rangle \leqslant \mathcal{S}_6$, and $S = \{2, 3, 5\}$. We intend to find the minimal image $\mathrm{Min}(G, S, \leqslant)$, where the ordering on subsets of Ω is the induced ordering from $\leqslant$ on Ω as explained in Definition 2.3.*

We split our problem into pieces by looking at cosets of $G_1 = \langle (3, 4, 5) \rangle$. The minimal image of S under G will be realized by an element contained in (at least) one of the cosets of G_1, so if we find the minimal image of S under

elements in each coset, and then take the minimum of these, we will find the global minimum.

Lemma 2.6 gives that for all $g \in G$ *it holds that* $\mathrm{Min}(g \cdot G_1, S, \leqslant) = \mathrm{Min}(G_1, S^g, \leqslant)$*, and so we can change our problem from looking for the minimal image of* S *with respect to cosets of* G_1 *to looking at images of* S^g *under elements of* G_1 *where* g *runs over a set of coset representatives of* G_1 *in* G*.*

For each $i \in \{1, ..., 6\}$ *we need an element* $g_i \in G_{i \mapsto 1}$ *(where any exist), so that we can then consider* S^{g_i}*.*

We choose the elements id, (162)*,* (146532)*,* $(14)(23)(56)$*,* (142365) *and* (126) *and obtain six images of* S*:*

$$\{2,3,5\}, \{1,3,5\}, \{3,1,2\}, \{3,2,6\}, \{3,6,1\}, \{6,3,5\}.$$

As we are looking at the images of these sets under G_1*, we know that all images of a set containing* 1 *will contain* 1*, and all images of a set not containing* 1 *will not contain* 1*. From Definition 2.3, all subsets of* $\{1, \ldots, 6\}$ *containing* 1 *are smaller than all subsets not containing* 1*. This means that we can filter our list down to* $\{1,3,5\}, \{3,1,2\}$ *and* $\{3,6,1\}$*.*

Furthermore, G_1 *fixes* 2*, so by the same argument we can filter our list of sets not containing* 2*, leaving only* $\{3,1,2\}$*. The minimal image of this under* G_1 *is clearly* $\{3,1,2\}$ *(in this particular case we could of course also have stopped as soon as we saw* $\{3,1,2\}$*, as this is the smallest possible set of size 3).*

Now, let us consider what would happen if the ordering of the integers was reversed, so we are looking for $\mathrm{Min}(G, S, \geqslant)$*, again with the induced ordering.*

For the same reasons as above, we begin by calculating $G_6 = \langle (3,5,4) \rangle$ *and by finding images of* S *for some element from each coset of* G_6 *in* G*.*

An example of six images is

$$\{1,5,4\}, \{6,5,3\}, \{4,6,1\}, \{5,1,2\}, \{3,2,6\}, \{2,3,5\}.$$

We can ignore anything that does not contain 6*, so we are left with:*

$$\{6,5,3\}, \{4,6,1\}, \{3,2,6\}.$$

As 5 *is not fixed by* G_6*, we can not reason about the presence or absence of* 5 *in our sets. There is an image of every set which contains* 5*, and there*

are even two distinct images of $\{6,5,3\}$ *which contain* 5*. Therefore we must continue our search by considering* $G_{6,5}$*.*

Application of an element from each coset of $G_{6,5}$ *to* S *generates nine sets, of which four contain the element* 5*. In fact we reach* $\{6,3,5\},\{6,5,4\}$ *from the set* $\{6,4,3\}$*, we reach* $\{5,6,1\}$ *from the set* $\{4,6,1\}$ *and we reach* $\{5,2,6\}$ *from the set* $\{3,2,6\}$*. From these we extract the minimal image* $\{6,5,4\}$*.*

In this example, different orderings of $\{1,2,3,4,5,6\}$ *produced different sized searches, with different numbers of levels of search required.*

3. Minimal Images Under Alternative Ω orderings

As was demonstrated in Example 2.7, the choice of ordering of the set our group acts on influences the size of the search for a minimal image. In this section we will show how to create orderings of Ω which, on average, reduce the size of search for a minimal image.

We begin by showing how large a difference different orderings can make. We do this by proving that, for any choice $\preccurlyeq$ of ordering of Ω, group G and any input set S, we can construct a minimal image problem that is as hard as finding $\mathrm{Min}(G,S,\preccurlyeq)$, but where reversing the ordering on Ω makes the problem trivial.

Given $n \in \mathbb{N}$, a permutation group G on $\{1,...,n\}$ with some ordering $\preccurlyeq$ and a subset $S \subseteq \{1,...,n\}$, we construct a group H and a set T such that $\mathrm{Min}(G,S,\preccurlyeq) = \mathrm{Min}(H,T,\preccurlyeq) \cap \{1..n\}$, which shows that finding $\mathrm{Min}(H,T,\preccurlyeq)$ is at least as hard as finding $\mathrm{Min}(G,S,\preccurlyeq)$. On the other hand, we will show that $\mathrm{Min}(H,T,\succcurlyeq) = T$, and therefore no search is required at all. This is done in Lemma 3.5. An example along the way will illustrate the construction.

Definition 3.1. *We fix* $n \in \mathbb{N}$ *and we let* $k \in \mathbb{N}$*. For all* $j \in \mathbb{N}$ *we define* $q(j) \in \mathbb{N}$ *and* $r(j) \in \{1,...,n\}$ *such that* $j = q(j) \cdot n + r(j)$*.*

Let $\mathrm{ext} : G \to \mathcal{S}_{k \cdot n}$ *be the following map: For all* $g \in G$ *and all* $j \in \{1,...,k \cdot n\}$*, the element* $\mathrm{ext}(g)$ *maps* j *to* $q(j) \cdot n + r(j)^g$*.*

Example 3.2. *Let* $n = 4$ *and* $G = \mathcal{S}_4$*. Then we extend the action of* G *to the set* $\{1,...,12\}$ *using the map* ext*.*

For example $g = (134)$ *maps* 4 *to* 1*. We write* $12 = 2 \cdot 4 + 4$ *and then it follows that* $\mathrm{ext}(g)$ *maps* 12 *to* $2 \cdot 4 + 4^g = 8 + 1 = 9$*. In fact* g *acts*

simultaneously on the three tuples $(1,2,3,4)$*,* $(5,6,7,8)$ *and* $(9,10,11,12)$ *as it does on* $(1,2,3,4)$*.*

Definition 3.3. *Fixing* $n,k \in \mathbb{N}$ *and a subgroup* G *of* $\mathcal{S}_n$*, and using the map* ext *defined above, we say that* H ***is the extension of*** G ***on*** $\{1,...,k\cdot n\}$ *if and only if* $H = \{\mathrm{ext}(g) \mid g \in G\}$ *is the image of* G *under the map* ext*.*

The extension H of G on a set $\{1,...,k\cdot n\}$ is a subset of $\mathcal{S}_{k\cdot n}$. We show now that even more is true:

Lemma 3.4. *Let* $n,k \in \mathbb{N}$ *and* $G \leqslant \mathcal{S}_n$*. Then the extension of* G *onto* $\{1,...,k\cdot n\}$ *is a subgroup of* $\mathcal{S}_{k\cdot n}$ *that is isomorphic to* G*.*

Proof. Let $H := \mathrm{ext}(G)$ be the image of G under the map ext and let $a,b \in G$ be distinct. Then let $j \in \{1,...,n\}$ be such that $j^a \neq j^b$. By definition $\mathrm{ext}(a)$ and $\mathrm{ext}(b)$ map j in the same way that a and b do, so we see that $\mathrm{ext}(a) \neq \mathrm{ext}(b)$. Hence the map ext is injective. Therefore $\mathrm{ext} : G \to H$ is bijective.

Next we let $a,b \in G$ be arbitrary and we let $j \in \{1,...,k\cdot n\}$. Then the composition ab is mapped to $\mathrm{ext}(ab)$ which maps j to $q(j)\cdot n + r(j)^{ab}$. Now $r(j)^{ab} = (r(j)^a)^b$ and therefore the composition $\mathrm{ext}(a)\,\mathrm{ext}(b) \in \mathcal{S}_{k\cdot n}$ maps j to $(q(j)\cdot n + r(j)^a)^{\mathrm{ext}(b)} = q(j)\cdot n + (r(j)^a)^b$. This is because $r(j)^a \in \{1,...,n\}$.

Hence $\mathrm{ext}(ab) = \mathrm{ext}(a)\,\mathrm{ext}(b)$, which means that ext is a group homomorphism and hence that G and its image are isomorphic. □

Lemma 3.5. *Let* $n \in \mathbb{N}$ *and* $G \leqslant \mathcal{S}_n$*. Let* H *denote the extension of* G *on* $\{1,...,(n+1)\cdot n\}$ *and let* $S \subseteq \{1,...,n\}$*. Let* $T := S \cup \{l\cdot n + l \mid l \in \{1..n\}\}$*,* $\leqslant$ *denote the natural ordering of the integers, and* $\geqslant$ *its reverse. For simplicity we use the same symbols for the ordering induced on* $\mathcal{P}(\Omega)$*, respectively. Then*

- $\mathrm{Min}(H,T,\leqslant) \cap \{1,\ldots,n\} = \mathrm{Min}(G,S,\leqslant)$.
- $\mathrm{Min}(H,T,\geqslant) = T$.

Proof. Let $h \in H$. Then by construction h stabilizes the partition $[1,...,n|n+1,...,2n|...|n\cdot n,...,(n+1)\cdot n]$. Moreover, for all $i \in \{1,...,n\}$ and $g \in G$ we have that $i^g = i^{\mathrm{ext}(g)}$ and so Lemma 3.4 implies that

$$\begin{aligned}\mathrm{Min}(G,S,\leqslant) &= \min_{\leqslant}\{S^g \mid g\in G\}\cap\{1,\ldots,n\}\\ &= \min_{\leqslant}\{S^{\mathrm{ext}(g)} \mid g\in G\}\cap\{1,\ldots,n\}\\ &= \min_{\leqslant}\{S^h \mid h\in H\}\cap\{1,\ldots,n\}\\ &= \mathrm{Min}(H,T,\leqslant)\cap\{1,\ldots,n\}.\end{aligned}$$

This proves the first statement.

For the second statement we notice that $(n+1)\cdot n$ is now the smallest element of T, and it cannot be mapped to anything smaller, because it also is the smallest element available. So if we let $h \in H$ be such that $T^h = \mathrm{Min}(H,T,\geqslant)$, then h fixes the point $n(n+1)$. By definition of the extension, it follows that h also fixes $k\cdot n$ for all $k\in\{1\ldots n\}$. The next point of T under the ordering is n^2-1. It cannot be mapped by h to n^2, because n^2 is already fixed, hence h has to fix n^2-1, too.

Arguing as above it follows that all points are fixed by h, thus in particular $\mathrm{Min}(H,T,\geqslant) = T$, as stated. □

3.1. Comparing Minimal Images Cheaply

We describe some important aspects of Linton's algorithm for computing the minimal image of a subset of Ω.

Definition 3.6. *Suppose that $(\Omega,\leqslant)$ is a totally ordered set and that $G\leqslant Sym(\Omega)$. Then $\mathrm{Orb}(G)$ denotes the list of orbits of G on Ω, ordered by the size of the smallest element in each orbit with respect to $\leqslant$.*

*A G-orbit will be called a **unit** if and only if it has size 1.*

*If $S\subseteq\Omega$, then we say that a G-orbit is **empty in S** if and only if it is disjoint from S as a set, and we say that it is **full in S** if and only if it is completely contained in S.*

Example 3.7. *Let $\Omega := \{1,...,8\}$, with the natural ordering on the integers, let $G := \langle(1,4),(2,8),(5,6),(7,8)\rangle$ and let $S := \{1,3,5,6\}$.*

Then $\mathrm{Orb}(G) = [\{1,4\},\{2,7,8\},\{3\},\{5,6\}]$ because this list contains all the G-orbits and they are ordered by the smallest element in each orbit,

namely 1 *in the first,* 2 *in the second,* 3 *in the third, which is a unit, and* 5 *in the last (because* 4 *is already in an earlier orbit).*

The orbits $\{3\}$ *and* $\{5,6\}$ *are full in* S*, the orbit* $\{2,7,8\}$ *is empty in* S *and* $\{1,4\}$ *is neither.*

Lemma 3.8. *Suppose that* $(\Omega, \leqslant)$ *is a totally ordered finite set and that* $G \leqslant Sym(\Omega)$*. If* $\mathrm{Min}(G, S, \leqslant) = \mathrm{Min}(G, T, \leqslant)$ *and* $\omega \in \Omega$*, then* ω^G *is empty in* S *if and only if it is empty in* T*, and* ω^G *is full in* S *if and only if it is full in* T*.*

Proof. Let $\omega \in \Omega$ and suppose that ω^G is empty in S. As ω^G is closed under the action of G, and $S_0 := \mathrm{Min}(G, S, \leqslant)$ is an image of S under the action of G, we see that ω^G is empty in S_0 and hence in $T_0 := \mathrm{Min}(G, T, \leqslant)$. Thus ω^G is empty in T, which is an image of T under the action of G. The same arguments work vice versa.

Next we suppose that ω^G is full in S. Then it is full in $S_0 = T_0$ and hence in T, and the same way we see the converse. □

We can now prove Theorem 3.9, which is the main technique used to reduce search. This allows us to prove the minimal image of some set S will be smaller than or equal to the minimal image of a set T, without explicitly calculating the minimal image of either S or T.

Theorem 3.9. *Suppose that* G *is a permutation group on a totally ordered finite set* $(\Omega, \leqslant)$ *with list of orbits* $O := \mathrm{Orb}(G)$*, and that* S *and* T *are two subsets of* Ω *where* $|S| = |T|$*.*

Suppose further that $o \in O$ *is the first orbit in the list* $\mathrm{Orb}(G)$ *that is neither full in both* S *and* T *nor empty in both* S *and* T*. If* o *is empty in* T*, but not in* S*, then* $\mathrm{Min}(G, S, \leqslant)$ *is strictly smaller than* $\mathrm{Min}(G, T, \leqslant)$*.*

Proof. Suppose that o is empty in T, but not in S. Then o is empty in $T_0 := \mathrm{Min}(G, T, \leqslant)$, but not in $S_0 := \mathrm{Min}(G, S, \leqslant)$, and in particular T_0 and S_0 are distinct, as we have seen in Lemma 3.8. If $\omega \in \Omega$, then the orbit ω^G is closed under the action of G. Thus, if $A \subseteq \Omega$ and ω^G is full in A, then for all $g \in G$ it is true that ω^G is full in A^g. Similarly, if ω^G is empty in A, then ω^G is empty in A^g for all $g \in G$.

Let $\alpha := \min_{\leqslant}(o)$ and let $\omega \in \Omega$. If $\omega \lneq \alpha$, then $\omega \notin o$, which means that the orbit ω^G appears in the list O *before* o does. Then the choice of o implies that one of the following two cases holds:

(i) ω^G is full in both S and T. In particular, for all $g \in G$ we have that $\omega \in S^g \cap T^g$.

(ii) ω^G is empty in of both S and T. In particular, for all $g \in G$ we have that $\omega^G \cap S^g = \varnothing$ and $\omega^G \cap T^g = \varnothing$.

If S_0 contains an element $\omega \in \Omega$ such that $\omega \lneq \alpha$, then Case (i) above holds and $\omega \in T_0$. So $S_0 \cap \{\omega' \in \Omega \mid \omega' \lneq x\} = T_0 \cap \{\omega' \in \Omega \mid \omega \lneq \alpha\}$.

Since S_0 and T_0 are distinct, they must differ amongst the elements at least as large as α and, since they have the same cardinality, the smallest such element determines which one of S_0 and T_0 is smaller.

We recall that $o = \alpha^G$ is empty in T and non-empty in S, so there exists some $g \in G$ such that $\alpha \in S^g$. Then $S^g = S_0$ and $\alpha \notin T_0$, so S_0 is strictly smaller that T_0. □

Here is an example how to use Theorem 3.9.

Example 3.10. *Let* $\Omega := \{1, ..., 10\}$ *with natural ordering, and let* $G := \langle (12), (45), (56), (89) \rangle$. *We consider the sets* $S := \{3, 6, 7\}$ *and* $T := \{3, 7, 9\}$ *and we want to calculate the smallest of* $Min(G, S)$ *and* $Min(G, T)$. *Hence, we want to know which one is smaller as cheaply as possible, to avoid superfluous calculations.*

We first list the orbits of G*:* $[\{1, 2\}, \{3\}, \{4, 5, 6\}, \{7\}, \{8, 9\}, \{10\}]$.

Going through the orbits as listed, we see that the first one is empty in S *and* T*, the second one is full in* S *and* T*, and the third one gives a difference for the first time. It is empty in* T*, but not in* S*, so Theorem 3.9 yields that the minimal image of* S *is strictly smaller than that of* T.

3.2. Static Orderings of Ω

In this section we look at which total ordering of Ω should be used to minimise the amount of time taken to find minimal images of subsets of Ω.

Given a group G we will choose an ordering on Ω such that orbits with few elements appear as early as possible. In particular, unit orbits should appear first.

This is justified by the fact that unit orbits are always either full or empty. Also, we would expect smaller orbits to be more likely to be empty or full than larger orbits. This means that small orbits placed early in the ordering of Ω are more likely to lead to Theorem 3.9 being applicable, leading to a reduction in search.

Algorithm 1 heuristically chooses a new ordering for an ordered set Ω, only depending on the group G, under the assumption that the algorithm that computes minimal images will pick a point from a smallest non-unit orbit to branch on. This will not always be true – in practice Linton's algorithm branches on the first orbit which contains some point contained in one of the current candidates for minimal image.

However, we will show that in Section 5 that Algorithm 1 produces substantially smaller, and therefore faster, searches in practice.

It is not necessary in Line 8 of Algorithm 1 to choose the smallest element of $Points$, choosing an arbitrary element will, on average, perform just as well. By fixing which point is chosen, we ensure that independent implementations will produce the same ordering and therefore the same canonical image.

Algorithm 1 FixedMinOrbit

```
 1: procedure MinOrbitOrder(Ω, G)
 2:     Remain := Ω
 3:     Order := []
 4:     H := G
 5:     while |Remain| > 0 do
 6:         OrbSize := Min {|o| | o ∈ Orb(H), o ∩ Remain ≠ ∅}
 7:         Points := {o | o ∈ Orb(H), |o| = OrbSize, o ∩ Remain ≠ ∅}
 8:         MinPoint := Min {x|o ∈ Points, x ∈ o}
 9:         Remain := Remain\{MinPoint}
10:         Add(Order, MinPoint)
11:         H := G_MinPoint
12:     return Order
```

We will also consider one simple modification of Algorithm 1, `FixedMaxOrbit`, which is the same as `FixedMinOrbit`, with line 6 changed to pick orbits of

maximum size.

If our intuition about Theorem 3.9 is correct, then `MaxOrbit` should almost always produce a larger search than `MinOrbit`, or a random ordering of Ω.

3.3. Implementing alternative orderings of Ω

Having calculated an alternative *Order* using **FixedMinOrbit** or **FixedMaxOrbit**, we could create a version of **MinimalImage** which accepted an explicit ordering. However, rather than editing the algorithm, we can instead perform a pre-processing step, using Lemma 3.11.

Lemma 3.11. *Consider a group G that acts on $\Omega = \{1, \ldots, n\}$ and a permutation $\sigma \in Sym_\Omega$. We define an ordering $\leqslant_\sigma$ on $\{1, \ldots, n\}$, where for all $x, y \in \Omega$ we have that $x \leqslant_\sigma y$ if and only if $x^\sigma \leqslant y^\sigma$.*

For the induced orderings $\preccurlyeq$ and $\preccurlyeq_\sigma$ on subsets of Ω as in Definition 2.3 it holds that

$$X \preccurlyeq_\sigma Y \Leftrightarrow X^\sigma \preccurlyeq Y^p$$

for all subsets X and Y of Ω, and hence (simplifying notation)

$$Min(G, S, \leqslant_\sigma) = Min(G^\sigma, S^\sigma, \leqslant)^{\sigma^{-1}}.$$

Proof. Following Definition 2.3, $X \preccurlyeq_\sigma Y$ if and only if there is an $x \in X$ such that $x \notin Y$ and for all $y \in Y \backslash X$ it holds that $x \leqslant_\sigma y$. By definition of $\leqslant_\sigma$, this is the case whenever $x^\sigma \leqslant y^\sigma$, and since $x^\sigma \in X^\sigma$ and for all y^σ in $Y^\sigma \backslash X^\sigma$ it holds that $x^\sigma \leqslant y^\sigma$, it follows that $X^\sigma \preccurlyeq Y^\sigma$.

Consider the map $\varphi_\sigma : S^G \to (S^\sigma)^{G^\sigma}$, which maps sets $X \in S^G$ to $X^\sigma \in (S^\sigma)^{G^\sigma}$. This map is bijective, and by the above it respects the order, so the second claim follows.

□

Lemma 3.11 gives an efficient method to calculate minimal images under different orderings without having to alter the underlying algorithm. The most expensive part of this algorithm is calculating G^σ, but this is still very efficiently implemented in systems such as GAP, and also can be cached so it only has to be calculated once for a given G and σ.

4. Dynamic Ordering of Ω

In Section 3.2, we looked at methods for choosing an ordering for Ω which allows a minimal image algorithm to search more quickly. There is a major limitation to this technique – it does not make use of the sets whose canonical image we wish to find.

In this section, instead of producing an ordering ahead of time, we will incrementally define the ordering of Ω as the algorithm progresses. At each stage we will consider exactly which extension of our partially constructed ordering will lead to the smallest increase in the number of sets we must consider.

We are not free to choose our ordering arbitrarily as we must still map two sets in the same orbit of G to the same canonical image. However, we can use different orderings for sets which are in different orbits of G.

Firstly, we will explain how we build the orderings which our algorithm uses.

4.1. Orderings

When building canonical images, we build orderings as the algorithm progresses. We represent these partially built orderings as ordered partitions.

Definition 4.1. *Given $k \in \mathbb{N}$ and $P = [X_1, ..., X_k]$, an ordered partition of $\mathcal{P}(\Omega)$, then given two subsets S and T of Ω, we define $S <_P T$ if and only if the cell which contains S occurs before the cell which contains T in P.*

*We say that P is G-*invariant *if and only if for all $i \in \{1, ..., k\}$ and $g \in G$ it holds that $S \in X_k$ if and only if $S^g \in X_k$.*

A refinement *of an ordered partition $[X_1, ..., X_k]$ is an ordered partition $[Y_{1,1}, Y_{1,2}, ..., Y_{k,l}]$ where $l \in \mathbb{N}$ and such that, for all $i \in \{1, ..., k\}$ and $j \in \{1, ..., l\}$, we have that $Y_{i,j} \subseteq X_i$.*

A completion *of an ordered partition $X = [X_1, ..., X_k]$ is a refinement where every cell is of size one. Given an ordering for Ω, the* standard completion *of an ordered partition X orders the members of each cell of X using the ordering on sets from Definition 2.3.*

In our algorithm we need a completion of an ordered partition, but the exact completion is unimportant – it is only important that given an ordered

partition X, we always return the same completion. For this reason we define the standard completion of an ordered partition.

Example 4.2. *Let* $G := \langle (12) \rangle \leqslant \mathcal{S}_3$ *and* $\Omega := \{1, 2, 3\}$. *Moreover let*

$P := [\{\}, \{1\}, \{2\}, \{1, 3\}, \{2, 3\} \mid \{3\}, \{1, 2\}, \{1, 2, 3\}]$ *be an ordered partition of* $\mathcal{P}(\Omega)$.

The orbits of G *on* Ω *are* $\{1, 2\}$ *and* $\{3\}$. *In particular all elements of* G *stabilize the partition* P.

The ordered partition $Q := [\{1, 3\}, \{2\} \mid \{\}, \{1\}, \{2, 3\} \mid \{3\}, \{1, 2\}, \{1, 2, 3\}]$ *is a refinement of* P *that is not* G*-invariant.*

To see this, we let $g := (12) \in G$. *We have that* $\{1, 3\}$ *is in the first cell of* Q, *but* $\{1, 3\}^g = \{2, 3\}$, *which is not in the first cell.*

In Example 4.2 we only considered a very small group, as the size of $\mathcal{P}(\Omega) = 2^{|\Omega|}$. In practice we will not explictly create ordered partitions of $\mathcal{P}(\Omega)$, but instead store a compact description of them, from which we can deduce the cell that any particular set is in.

In this paper, we will consider two methods of building and refining ordered partitions. We first define the orbit count of a set, which we will use when building refiners.

Definition 4.3. *Let* G *be a group acting on an ordered set* Ω, *and* $S \subseteq \Omega$. *Define the* ***orbit count*** *of* S *in* G, *denoted* $\mathrm{Orbcount}(G, S)$ *as follows: Given the list* $\mathrm{Orb}(G)$ *of orbits of* G *on* Ω *sorted by the their smallest member, the list* $\mathrm{Orbcount}(G, S)$ *contains the size of the intersection* $|o \cap S|$ *in place of* $o \in \mathrm{Orb}(G)$.

We will see the practical use of Orbcount in Lemma 5.1.

Lemma 4.4. *Suppose that* $(\Omega, \leqslant)$ *is a totally ordered finite set and that* $G \leqslant Sym(\Omega)$. *Then given two sets* $S, T \subseteq \Omega$ *such that there is* $g \in G$ *such that* $S^g = T$, *then* $\mathrm{Orbcount}(G, S) = \mathrm{Orbcount}(G, T)$.

Proof. Let $o \in \mathrm{Orb}(G)$ and $g \in G$ with $S^g = T$. Then $o^g = o$ and $\alpha \in (o \cap S)$ if and only if $\alpha^g \in (o \cap S)^g$, if and only if $\alpha^g \in (o^g \cap S^g) = (o \cap T)$.

□

Definition 4.5. *Let* P *be an ordered partition of* $\mathcal{P}(\Omega)$.

- *If $\alpha \in \Omega$, then the **point refinement** of P by α is the ordered partition Q, where each cell X_i of P is split into two cells, namely $\{S \mid S \in X_i, \alpha \in S\}$, and $\{S \mid S \in X_i, \alpha \notin S\}$. If one of these sets is empty, X_i is not split.*

- *If $G \leqslant Sym(\Omega)$ and $C = \mathrm{Orbcount}(G, T)$ for some set $T \subseteq \Omega$, the **orbit refinement** of P by C is the ordered partition Q, where each cell X_i of P is split into two cells, namely $\{S \mid S \in X_i, \mathrm{Orbcount}(G, S) = C\}$, and $\{S \mid S \in X_i, \mathrm{Orbcount}(G, S) \neq C\}$. If one of these sets is empty, X_i is not split.*

4.2. Algorithm

We will now present our algorithm. First, we give a technical definition which will be used in proving our algorithm correct.

Definition 4.6. *For all $n \in \mathbb{N}$ we define $\mathcal{L}_n$ to be the set of lists of length n whose entries are non-empty subsets of Ω. If $X \in \mathcal{L}_n$, then as a convention we write $X_1, ..., X_n$ for the entries of the list X.*

If $X \in \mathcal{L}_n$ and $H \leqslant G \leqslant \mathcal{S}_n$ is such that $|G : H| = k \in \mathbb{N}$ and $Q = \{q_1, ..., q_k\}$ is a set of coset representatives of H in G, then we define X^Q to be the list whose first k entries are $X_1^{q_1}, ..., X_1^{q_k}$, followed by $X_2^{q_1}, ..., X_2^{q_k}$ until the last k entries are $X_n^{q_1}, ..., X_n^{q_k}$. We note that $X^Q \in \mathcal{L}_{n \cdot k}$.

Let $X, Y \in \mathcal{L}_n$. We say that X and Y are G-equivalent if and only if there exist a permutation σ of $\{1, ..., n\}$ and group elements $g_1, ..., g_n \in G$ such that, for all $i \in \{1, ..., n\}$ it holds that $Y_i = X_{i^\sigma}^{g_i}$.

We now prove a series of three lemmas about coset representatives, which form the basis for the correctness proof of our algorithm. They are used to perform the recursive step, moving from a group to a subgroup.

Lemma 4.7. *Suppose that G is a permutation group on a set Ω, that H is a subgroup of G of index $k \in \mathbb{N}$ and that T is a set of left coset representatives of H in G.*

Then the following is true:

(i) $|T| = k$.

(ii) *If* $T = \{t_1, ..., t_k\}$ *and* $g \in G$ *and if, for all* $i \in \{1, ..., k\}$ *we define* $q_i := gt_i$*, then* $Q := \{q_1, ..., q_k\}$ *is also a set of coset representatives of* H *in* G*. In particular there is a bijection from* Q *to any set of left coset representatives of* H *in* G*.*

Proof. By definition the index of H in G is the number of (left or right) cosets of H in G.

For the second statement we let $i, j \in \{1, ..., k\}$ be such that $q_iH = q_jH$, hence $gt_iH = gt_jH$. Then $t_j^{-1}t_i = t_j^{-1}g^{-1}gt_i \in H$ and hence $t_iH = t_jH$. Hence $i = j$ because t_i and t_j are from a set of coset representatives. □

Lemma 4.8. *Suppose that* G *is a permutation group on a set* Ω *and that* $S, T \subseteq \Omega$ *are such that the lists* $[S]$ *and* $[T]$ *are* G*-equivalent.*

Let H *be a subgroup of* G *of index* $k \in \mathbb{N}$ *and let* $P = \{p_1, ..., p_k\}$ *and* $Q := \{q_1, ..., q_k\}$ *be sets of left coset representatives of* H *in* G*.*

Then $[S]^P$ *and* $[T]^Q$ *are* H*-equivalent.*

Proof. As $[S]$ and $[T]$ are G-equivalent, we know that there exists a group element $g \in G$ such that $S^g = T$.

We fix g, for all $i \in \{1, ..., k\}$ we let $t_i := gq_i$ and we consider the set $T := \{t_i | i \in \{1, \ldots, k\}\}$. Then T is also a set of left coset representatives of H in G, by Lemma 4.7. As P is also a set of left coset representatives, we know that T and P have the same size, so there is a bijection from P to T. This can be expressed in the following way:

There is a permutation $\sigma \in \mathcal{S}_k$ such that, for all $i \in \{1, .., k\}$, it is true that $p_{i^\sigma}H = t_iH$, which means that there is a unique $h_i \in H$ such that $p_{i^\sigma}h_i = t_i$.

Let now $S_i := S^{p_i}$ and $T_i := T^{q_i}$ for $i \in \{1, ..., k\}$, then

$$T_i = T^{q_i} = (S^g)^{q_i} = S^{t_i} = S^{p_{i^\sigma}h_i} = (S^{p_{i^\sigma}})^{h_i} = (S_{i^\sigma})^{h_i},$$

hence $[S]^P$ and $[T]^Q$ are H-equivalent. □

Lemma 4.9. *Suppose that* G *is a permutation group on a set* Ω*, that* $n \in \mathbb{N}$ *and that* $X, Y \in \mathcal{L}_n$ *are* G*-equivalent. Let* H *be a subgroup of* G *of index* $k \in \mathbb{N}$ *and let* $P = \{p_1, ..., p_k\}$ *and* $Q := \{q_1, ..., q_k\}$ *be sets of left coset representatives of* H *in* G*.*

Then X^P *and* Y^Q *are* H*-equivalent.*

Proof. As X and Y are G-equivalent, we know that there exist a permutation $\sigma \in \mathcal{S}_n$ and $g_1, ..., g_n \in G$ such that, for all $i \in \{1, ..., n\}$: $Y_i = X_{i^\sigma}^{g_i}$. We fix this permutation σ.

If $i \in \{1, \ldots, n\}$, then $[X_{i^\sigma}]$ and $[Y_i]$ satisfy the hypothesis of Lemma 4.8, so it follows that $[X_{i^\sigma}]^P$ and $[Y_i]^Q$ are H-equivalent.

So there exist a permutation $\alpha_i \in \mathcal{S}_k$ and group elements $h_{i1}, ..., h_{ik} \in H$ such that $(X_{i^\sigma}^{p_j{}^{\alpha_i}})^{h_{ij}} = Y_i^{q_i}$ for all $j \in \{1, ..., k\}$.

Using σ and $\alpha_1, ..., \alpha_k$ we define a permutation γ on $1, ..., n \cdot k$:

We express $l \in \{1, ..., n \cdot k\}$ uniquely as $l = c_l \cdot k + r_l$ where $c_l \in \{0, ..., n-1\}$ and $r_l \in \{1, ..., k\}$ and we define

$$l^\gamma := (c_l + 1)^\sigma \cdot k + r_l^{\alpha_{c_l+1}}.$$

This is well-defined because of the ranges of c_l and r_l and it is a permutation because of the uniqueness of the expression and because σ and $\alpha_1, ..., \alpha_n$ are permutations.

For each $l \in \{1, ..., n \cdot k\}$, expressed as $l = c_l \cdot k + r_l$ as we did above, we set $h := h_{c_l+1, r_l}$, $X'_l := X_{c_l+1}^{p_{r_l}}$ and $Y'_l := Y_{c_l+1}^{q_{r_l}}$.

Then $X^P = [X'_1, ..., X'_{n \cdot k}]$ and $Y^Q = [Y'_1, ..., Y'_{n \cdot k}]$.

If we set $p := p_{r_l{}^{a_{c_l+1}}}$ and $q := q_{r_l}$, then we have, for all $l = c_l \cdot k + r_l$, that

$$X'^h_{l^\gamma} = (X^p_{(c_l+1)^\sigma})^h = ((X^\sigma_{c_l+1})^p)^h = Y^q_{c_l+1} = Y'_l.$$

This is H-equivalence. □

We can now describe the algorithm we use to compute canonical images, and prove that it works correctly.

Definition 4.10. *Suppose that Ω is a finite set, that G is a permutation group on Ω, that $L \in \mathcal{L}_k$ and that P is an ordered partition on $\mathcal{P}(\Omega)$.*

- *An* Ω-selector *is a function $\mathcal{S}$ such that*
  - *$\mathcal{S}(\Omega, G, L, P) = \omega \in \Omega$, where $|\omega^G| > 1$;*
  - *$\mathcal{S}(\Omega, G, L, P) = \mathcal{S}(\Omega, G, M, P)$ whenever L and M from $\mathcal{L}_k$ are G-equivalent.*

- *An* Ordering refiner *is a function $\mathcal{O}$ such that for all G-invariant partitions P of $\mathcal{P}(\Omega)$*
  - *$\mathcal{O}(\Omega, G, L, P) = P'$, where P' is a G-invariant refinement of P;*
  - *$\mathcal{O}(\Omega, G, L, P) = \mathcal{O}(\Omega, G, M, P)$ whenever L and M from $\mathcal{L}_k$ are G-equivalent.*

An ordering refiner cannot return a total ordering, unless G acts trivially, as the partial ordering cannot distinguish between values contained in orbits of G.

Our method for finding canonical images is outlined in Algorithm 2. It recursively searches for the minimal image of a collection of lists, refining the ordering that is used as search progresses.

Algorithm 2 CanImage

Require: $\mathcal{S}$ is an Ω-selector, $\mathcal{O}$ is an ordering refiner

1: **procedure** CanImageRecurse(Ω, G, L, P)
2: **if** $|G| = 1$ **then**
3: $P' :=$ Standard completion of P
4: **return** Smallest member of L under P'
5: $H := G_{\mathcal{S}(\Omega,G,L,P)}$
6: $Q :=$ coset representatives of H in G
7: $P' := \mathcal{O}(\Omega, H, L^Q, P)$
8: $L' := [S | S \in L^Q, \nexists T \in L^Q . T <_{P'} S]$
9: **return** CanImageRecurse(Ω, H, L', P')

procedure CanImageBase($\Omega, InputG, InputS$)
2: **return** CanImageRecurse($\Omega, InputG, [InputS], [\mathcal{P}(\Omega)]$)

Theorem 4.11. *Suppose that Ω is a finite set, that G is a permutation group on Ω, and that $X \subseteq \Omega$. Then* CanImageBase$(\Omega, G, X) \in X^G$.

Proof. In every step of Algorithm 2 the list of considered sets is a list of elements of X^G. □

Theorem 4.12. *Let Ω be a finite set and G a permutation group on Ω. Let $X, Y \in \mathcal{L}_k$ be G-equivalent, and let P be a G-invariant ordered partition of $\mathcal{P}(\Omega)$. Then*

$$\textsc{CanImage}(\Omega, G, X, P) = \textsc{CanImage}(\Omega, G, Y, P).$$

Proof. We proceed by induction on the size of G.

The base case is $|G| = 1$. As X and Y are G-equivalent, X and Y contain the same sets, possibly in a different order. For a given P and Ω, there is only one standard completion of P which gives a complete ordering on $\mathcal{P}(\Omega)$, and so X and Y have the same smallest element under the standard completion of P and so the claim follows.

Consider now any non-trivial group G, and assume that the claim holds for all groups H where $|H| < |G|$.

Definition 4.10 and the fact that $\mathcal{S}$ is an Ω-selector imply that

$$H = G_{\mathcal{S}(\Omega,G,X,P)} = G_{\mathcal{S}(\Omega,G,Y,P)}.$$

Moreover, H is a proper subgroup of G. We take two sets Q_1 and Q_2 of coset representatives of H in G, which are not necessarily equal.

Since $\mathcal{O}$ is an ordering refiner, it holds that

$$P' = \mathcal{O}(\Omega, H, X, P) = \mathcal{O}(\Omega, H, Y, P)$$

By Lemma 4.9, X^{Q_1} and Y^{Q_2} are H-equivalent, and by definition P' is G-invariant. If we identify the cell P'_i of P' which contains the smallest element of X^{Q_1}, then L'_X contains those elements of X^{Q_1} which are in P'_i. Each of these elements is H-equivalent to an element of Y^{Q_2}, and therefore L'_X is H-equivalent to L'_Y. Then the induction assumption yields that

$$\textsc{CanImage}(\Omega, H, X', P) = \textsc{CanImage}(\Omega, G, Y', P),$$

so the claim follows by induction. □

We see from Theorem 4.11 and 4.12 that

$$\textsc{CanImage}(\Omega, G, X, P) = \textsc{CanImage}(\Omega, G, Y, P) \text{ if and only if } Y \in X^G.$$

We note that Algorithm 2 can easily be adapted to return an element g of G such that $X^g = \textsc{CanImage}(\Omega, G, X, P)$, by attaching to each set, when it is created, the permutation which maps it to the original input S. We omit this addition for readability.

5. Experiments

In this section, we will compare how well our new algorithms perform in comparison with the **MinimalImage** function of Linton's.

All of our experiments are performed in GAP [3] using code available in the `Images` package [7]. Where our algorithm requires it, we use the implementation of partition backtracking provided in the `Ferret` package [6].

We consider a selection of different canonical image algorithms and we analyze how they perform compared to each other, and compared to the traditional minimal image algorithm of Linton's, which we will refer to as **MinImage**.

The first three algorithms that we consider come from Section 3. They produce, given a group G on a set Ω, an ordering of Ω. This ordering is then used in **MinImage**.

(i) **FixedMinOrbit** uses results from Section 3 to calculate an alternative ordering of Ω, choosing small orbits first.
(ii) **FixedMaxOrbit** works similarly to **FixedMinOrbit**, choosing large orbits first.

We also consider algorithms which dynamically choose which value to branch on as search progresses. We will use the following lemma for the proof of correctness for all our orderings.

Lemma 5.1. *Let Ω be a finite set, G a permutation group on Ω, and P an ordered partition of $\mathcal{P}(\Omega)$. Let $L \in \mathcal{L}_n$, and $Count = [\mathrm{Orbcount}(G, S) | S \in L]$.*

- *Any function which accepts (Ω, G, L, P) and returns some $\omega \in \Omega$ with $|\omega^G| > 1$, which is invariant under reordering the elements of $Count$, is an Ω-selector.*
- *Any function which accepts (Ω, G, L, P) and returns either P or the point refinement of P by some ω^G for $\omega \in \Omega$, and is invariant under permutation of the elements of $Count$, is an ordering refiner.*
- *Any function which accepts (Ω, G, L, P) and returns either P or the orbit refiner of P by some member of $Count$ and is invariant under permutation of the elements of $Count$, is an ordering refiner.*

Proof. The only thing we have to show is that, if L and M are G-equivalent, then any of the functions above yield the same result for inputs (Ω, G, L, P) and (Ω, G, M, P). By Lemma 4.4, any G-equivalent lists L and M will produce the same list $Count$, up to reordering of elements, and hence the claim follows. □

Firstly we define a list of orderings. Each of these orderings chooses an orbit, or list of orbits, to branch on – we will then make an Ω-selector by choosing the smallest element in any of the orbits selected, to break ties (we could choose any point, as long as we picked it consistently). Each of these algorithms operates on a list $L \in \mathcal{L}_k$. In each case we look for an orbit, ignoring orbits of size one (as fixing a point which was already fixed leads to the same group).

Firstly we will consider two algorithms which only consider the group, and not L:

(i) **MinOrbit** Choose a point from a shortest non-trivial orbit which has a non-empty intersection with at least one element of L.

(ii) **MaxOrbit** Choose a point from a longest non-trivial orbit which has a non-empty intersection with at least one element of L.

We also consider four algorithms which consider both the group, and L. In the following, for an orbit o

(i) **RareOrbit** minimises $\sum_{s \in L} |s \cap o|$,

(ii) **CommonOrbit** maximises $\sum_{s \in L} |s \cap o|$,

(iii) **RareRatioOrbit** minimises $\log(\sum_{s \in L} |s \cap o|)/|o|$,

(iv) **CommonRatioOrbit** maximises $\log(\sum_{s \in L} |s \cap o|)/|o|$.

The motivation for **RareOrbit** is that this is the branch which will lead to the smallest size of the next level of search – this exactly estimates the size of the next level if our ordering refiner only fixed a point in orbit o. We therefore expect, conversely, **CommonOrbit** to perform badly and to produce very large searches.

One limitation of **RareOrbit** is that it will favour smaller orbits – in general we want to minimize the size of the whole search. The idea here is that if we have two levels of search where we split on an orbit of size two, and each time create 10 times more sets is equivalent to splitting once on an orbit of size four, and creating 100 times more sets. **RareRatioOrbit** compensates for this. We expect that **CommonRatioOrbit** is the inverse of this, so we also expect it to perform badly.

For each of these orderings, we use the ordering refiner which takes each fixed point of G in their order in Ω, and performs a point refinement by each recursively in turn. By repeated application of Lemma 4.4, this is a G-invariant ordering refiner.

We also have a set of orderings which make use of orbit counting. To keep the number of experiments under control, in each case we used the **RareOrbit** strategy to choose which point to branch on next, but then also build an orbit refiner.

Given an unordered list of orbit counts,

(i) **RareOrbitPlusMin** chooses the lexicographically smallest one.
(ii) **RareOrbitPlusRare** chooses the least frequently occurring orbit count list (using the lexicographically smallest to break ties).
(iii) **RareOrbitPlusCommon** chooses the most frequently occurring orbit count list (using the lexicographically smallest to break ties).

5.1. Experiments

In this section we perform a practical comparison of our algorithms, and the **Minimum** algorithm of Linton, for three different families of problems: grid groups, m-sets, and primitive groups.

5.2. Experimental Design

We will consider three sets of benchmarks in our testing. In each experiment, given a group G of permutations of $\{1, \ldots, n\}$, we will run an experiment with each of our orderings to find the canonical image of a set of size $\left\lfloor \frac{n}{2} \right\rfloor$, $\left\lfloor \frac{n}{4} \right\rfloor$ and $\left\lfloor \frac{n}{8} \right\rfloor$.

We run our algorithms on a randomly chosen conjugate of each primitive group, to randomise the initial ordering of the integers the group is defined

over. The same conjugate is used of each group in all experiments, and when choosing a random subset of size x from a set S, we always choose the same random subset. We use a timeout of five minutes for each experiment. We force GAP to build a stabilizer chain for each of our groups, before we begin our algorithm, as this can in some cases take a long time.

For each size of set, and each ordering, we measure three things. The total number of problems solved, the total time taken to solve all problems, counting timeouts as 5 minutes, and the number of moved points of the largest group solved. Our experiments were all performed on an Intel(R) Xeon(R) CPU E5-2640 v4 running at 2.40GHz, with twenty cores. Each copy of GAP was allowed a maximum of 6GB of RAM.

5.2.1. Grid Groups

In this experiment, we look for canonical images of sets in Grid groups (defined below).

Definition 5.2. *Let $n \in \mathbb{N}$. The direct product $\mathcal{S}_n \times \mathcal{S}_n$ acts on the set $\{1, \ldots, n\} \times \{1, \ldots, n\}$ of pairs in the following way:*

For all $(i, j) \in \{1, \ldots, n\} \times \{1, \ldots, n\}$ and all $(\sigma, \tau) \in \mathcal{S}_n \times \mathcal{S}_n$ we define

$$(i, j)^{(\sigma,\tau)} := (i^\sigma, j^\tau).$$

*The subgroup $G \leqslant Sym(\{1, \ldots, n\} \times \{1, \ldots, n\})$ defined by this action is called the $n \times n$ **grid group**.*

We note that, while the construction of the grid group is done by starting with an n by n grid of points and permuting rows or columns independently of each other, we actually represent this group as a subgroup of $\mathcal{S}_{n \cdot n}$, and we do not assume prior knowledge of the grid structure of the action.

We ran experiments on the grid groups for grids of size 5×5 to 100×100, as described in section 5.2. The results of this experiment are given in Table 5.2.3.

The basic algorithm, **Minimum**, is only able to solve 22 problems within the timeout. **FixedMinOrbit** solves 43 problems, while being implemented as a simple pre-processing step to **Minimum**. The dynamic **MinOrbit** is

able to solve 55 problems, and the best orbit-based strategy, **SingleMaxOrbit**, solves 70 problems. However the advanced techniques, which filter by orbit lists, perform much better. Even ordering by the most common orbit list leads to solving over 185 problems, and the best strategy, **RareOrbitPlusRare**, solves 235 out of the total 288 problems.

Furthermore, for these large groups, the algorithms are still performing very small searches: For example **FixedMinOrbit**, on its largest solved problem with size $\lfloor \frac{n}{2} \rfloor$ sets and a grid of size(12×12), generates $793,124$ search nodes, while **RareOrbitPlusMin** produces only $183,579$ search nodes on its largest solved problem with $\lfloor \frac{n}{2} \rfloor$ sets (65×65).

5.2.2. M-Sets

Linton [10] considers, given integers n, m and k, taking the action of $\mathcal{S}_n$ on the set of all m-sets, and then finding the minimal image of a random set of k points.

We ran experiments for $n \in \{10, 15, \ldots, 100\}$ and $m = 2$, for $n \in \{10, 15, \ldots, 35\}$ and $m = 4$, for $n \in \{10, 15, 20\}$ and $m = 6$ and also in the case where $n = 10, m = 8$ as described in Section 5.2. We choose these 30 experiments as these were the problems on which any of our techniques were able to solve in under 5 minutes. The results of our experiments are shown in Table 5.2.3.

Similarly to our experiments on grid groups, we find that the standard **Minimum** algorithm is only able to solve a very small set of benchmarks. Some of the better algorithms, including **FixedMinOrbit**, are able to solve 9 problems. In particular, once again **MinOrbit** is not significantly better than **FixedMinOrbit**, although it is slightly faster on average over all problems.

However, the orbit-based strategies do much better, solving all the problems which we set. In the case of sets of size $\lfloor \frac{n}{8} \rfloor$, the best technique is able to solve all problems any technique can solve, in under 5 minutes. The largest solved problem, which was instance $n = 35, m = 4$ for a set of size $\lfloor \frac{n}{8} \rfloor$, is solved in only $6,594$ search nodes by **RareOrbitPlusMin**, while the largest solved problem of **Minimum**, $n = 15, m = 4$ takes $631,144$ search nodes.

5.2.3. Primitive Groups

In this experiment we look for canonical images of sets under the action of a primitive group which moves between 2 and $1,000$ points. We remove the

	Stab	$\lfloor \frac{n}{2} \rfloor$			$\lfloor \frac{n}{4} \rfloor$			$\lfloor \frac{n}{8} \rfloor$		
	Search	# solved	largest	time	# solved	largest	time	# solved	largest	time
RareOrbitPlusRare	F	56	4,225	12,549	88	8,464	10,474	91	9,216	13,663
RareOrbitPlusMin	F	54	4,225	13,207	89	9,604	10,546	90	9,025	13,484
RareOrbitPlusCommon	F	30	2,209	19,306	65	5,476	15,408	90	9,216	14,009
SingleMaxOrbit	F	11	225	24,671	27	961	25,795	52	3,136	23,532
RareOrbit	F	8	144	24,763	17	625	27,213	34	4,096	25,979
MinOrbit	F	8	144	24,739	16	400	27,735	31	1,521	26,706
CommonRatioOrbit	F	7	121	24,990	15	400	28,255	29	1,296	27,322
FixedMinOrbit	T	8	144	24,796	13	361	28,619	22	841	28,841
FixedMinOrbit	F	8	144	24,786	13	361	28,573	21	841	28,859
Minimum	F	4	64	25,819	8	144	30,039	10	196	31,851
Minimum	T	4	64	25,822	7	144	29,962	10	196	31,829
RareRatioOrbit	F	3	49	26,126	5	81	30,379	8	144	32,804
MaxOrbit	F	3	49	26,123	5	81	30,373	8	144	32,771
FixedMaxOrbit	F	3	49	26,120	5	81	30,374	8	144	32,725
FixedMaxOrbit	T	3	49	26,113	5	81	30,381	8	144	32,450
CommonOrbit	F	3	49	26,131	5	81	30,396	8	144	32,826

Table 1: Finding canonical images in grid groups

	Stab	$\lfloor \frac{n}{2} \rfloor$			$\lfloor \frac{n}{4} \rfloor$			$\lfloor \frac{n}{8} \rfloor$		
	Search	# solved	largest	time	# solved	largest	time	# solved	largest	time
RareOrbitPlusMin	F	28	27,405	1,290	30	52,360	835	30	52,360	251
RareOrbitPlusRare	F	27	12,650	1,235	30	52,360	1,107	30	52,360	322
RareOrbitPlusCommon	F	27	12,650	1,819	30	52,360	1,189	30	52,360	335
MinOrbit	F	9	1,365	6,701	22	6,435	4,143	28	27,405	1,203
RareOrbit	F	9	1,365	6,769	21	6,435	4,190	28	27,405	1,276
FixedMinOrbit	F	9	1,365	6,671	22	6,435	4,292	27	12,650	1,765
FixedMinOrbit	T	9	1,365	6,678	22	6,435	4,291	27	12,650	1,766
CommonRatioOrbit	F	9	1,365	6,782	21	6,435	4,215	28	27,405	1,348
Minimum	F	4	210	7,866	5	210	7,562	7	1,365	7,033
Minimum	T	4	210	7,843	5	210	7,573	7	1,365	7,037
SingleMaxOrbit	F	4	210	7,813	5	210	7,554	6	1,365	7,325
FixedMaxOrbit	T	4	210	7,853	5	210	7,691	6	210	7,569
RareRatioOrbit	F	4	210	7,923	5	210	7,642	5	210	7,500
MaxOrbit	F	4	210	7,915	5	210	7,695	5	210	7,500
FixedMaxOrbit	F	4	210	7,890	5	210	7,666	5	210	7,500
CommonOrbit	F	4	210	7,947	5	210	7,724	5	210	7,500

Table 2: Finding canonical images in M-Set groups

natural alternating and symmetric groups, as finding minimal and canonical images in these groups is trivial and can be easily special-cased. So we look at a total number of $5,845$ groups, each of which was successfully treated by at least one algorithm.

We perform the experiment as described in section 5.2. The results are given in Table 6.

All algorithms are able to solve a large number of problems. This is unsurprising as many primitive groups are quite small (for example the cyclic groups), so any technique is able to brute-force search many problems. However, we can still see that, for the hardest problems $\lfloor \frac{n}{8} \rfloor$, many algorithms outperform **Minimum**, and the techniques which use extra orbit-counting filtering solve 300 more problems, and they run much faster.

For the easiest set of problems, $\lfloor \frac{n}{8} \rfloor$, we see that the algorithm **RareOrbitPlusMin**, which usually performs best, solves slightly fewer problems. This is because there are a small number of groups where the extra filtering provides no search reduction, but still requires a small overhead in time. However, the total time taken is still much smaller, and the algorithm only fails to solve five problems. These five problems involve groups that are isomorphic to the affine general linear groups $\mathrm{AGL}(8,2)$, $\mathrm{AGL}(6,3)$ and $\mathrm{AGL}(9,2)$, and the projective linear group $\mathrm{PSL}(9,2)$. This suggests that the linear groups may be a source of hard problems for canonical image algorithms in the future.

5.2.4. Experimental Conclusions

Our experiments show that using **FixedMinOrbit** is almost always superior to **Minimum**. As implementing **FixedMinOrbit** requires a fairly small amount of code and time over **Minimum**, this suggests any implementations of Linton's algorithm should have **FixedMinOrbit** added, as this provides a substantial performance boost, for relatively little extra coding.

Algorithms which dynamically order the underlying set, such as **MinOrbit** and **RareOrbit** provide only a small benefit over **FixedMinOrbit**. Algorithms which add orbit counting provide a much bigger gain, often allowing solving problems on groups many orders of magnitude larger than before, thereby greatly advancing the state of the art.

	Stab	$\lfloor \frac{n}{2} \rfloor$		$\lfloor \frac{n}{4} \rfloor$		$\lfloor \frac{n}{8} \rfloor$	
	Search	# solved	time	# solved	time	# solved	time
RareOrbitPlusMin	F	5,689	64,983	5,749	39,564	5,840	10,287
RareOrbitPlusRare	F	5,656	81,814	5,738	43,202	5,825	19,025
RareOrbitPlusCommon	F	5,561	113,011	5,721	59,740	5,816	25,392
FixedMinOrbit	T	5,360	220,076	5,623	82,684	5,817	18,295
FixedMinOrbit	F	5,354	251,786	5,628	99,253	5,816	30,193
MinOrbit	F	5,348	263,053	5,641	104,241	5,844	27,720
RareOrbit	F	5,324	272,631	5,632	105,537	5,844	29,365
CommonRatioOrbit	F	5,323	277,050	5,629	107,561	5,844	32,123
SingleMaxOrbit	F	4,811	465,908	5,250	253,467	5,648	112,147
Minimum	T	4,723	390,334	5,163	242,477	5,631	88,048
Minimum	F	4,710	501,952	5,180	280,686	5,633	106,983
FixedMaxOrbit	F	4,659	514,618	5,119	298,321	5,587	132,508
FixedMaxOrbit	T	4,674	392,753	5,095	251,001	5,583	107,433
MaxOrbit	F	4,641	544,222	5,104	310,402	5,587	144,182
RareRatioOrbit	F	4,614	559,690	5,090	316,609	5,586	152,201
CommonOrbit	F	4,604	569,047	5,086	319,288	5,586	154,142

Table 3: Finding canonical images in Primitive groups

6. Conclusions

We present a general framework and new set of algorithms which find the canonical image of a set under the action of a permutation group. Our experiments show that our new algorithms outperform the previous state of the art, often by orders of magnitude.

Out basic framework runs on the concept of refiners and selectors and is not limited to finding only canonical images of subsets of Ω. In future work we will investigate families of refiners and selectors which allow finding canonical images for many other combinatorial objects.

Acknowledgements.

All authors thank the DFG (**Wa 3089/6-1**) and the EPSRC CCP CoDiMa (**EP/M022641/1**) for supporting this work. The first author would like to thank the Royal Society, and the EPSRC (**EP/M003728/1**). The second author would like to acknowledge support from the OpenDreamKit Horizon 2020 European Research Infrastructures Project (#676541). The first two authors thank the Algebra und Geometrie group at Martin-Luther Universität Halle-Wittenberg for their hospitality and inspiring environment. The third author wishes to thank the Computer Science Department of the University of St Andrews for its hospitality during numerous visits.

References

[1] P. T. Darga, K. A. Sakallah, and I. L. Markov. Faster symmetry discovery using sparsity of symmetries. In *2008 45th ACM/IEEE Design Automation Conference*, pages 149–154, June 2008.

[2] Andreas Distler, Chris Jefferson, Tom Kelsey, and Lars Kotthoff. *The Semigroups of Order 10*, pages 883–899. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.

[3] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.8.4*, 2016.

[4] Ian P Gent and Barbara M Smith. Symmetry breaking in constraint programming. In *Proceedings of the 14th European conference on artificial intelligence*, pages 599–603. IOS press, 2000.

[5] Derek F. Holt, Bettina Eick, and Eamonn A. O'Brien. *Handbook of Computational Group Theory*. Discrete Mathematics and its Applications. Chapman & Hall/CRC, Boca Raton, 2005.

[6] Christopher Jefferson. *Ferret - GAP Package*, 2016.

[7] Christopher Jefferson and Eliza Jonauskyte. *Images - GAP Package*, 2016.

[8] Tommi Junttila and Petteri Kaski. Engineering an efficient canonical labeling tool for large and sparse graphs. In David Applegate, Gerth Stølting Brodal, Daniel Panario, and Robert Sedgewick, editors, *Proceedings of the Ninth Workshop on Algorithm Engineering and Experiments and the Fourth Workshop on Analytic Algorithms and Combinatorics*, pages 135–149. SIAM, 2007.

[9] Jeffrey S. Leon. Permutation group algorithms based on partitions, I: Theory and algorithms. *Journal of Symbolic Computation*, 12(4):533–583, 1991.

[10] Steve Linton. Finding the smallest image of a set. In *Proceedings of the 2004 International Symposium on Symbolic and Algebraic Computation*, ISSAC '04, pages 229–234, New York, NY, USA, 2004. ACM.

[11] Brendan D. McKay. Practical graph isomorphism. *Congr. Numer.*, (30):45–87, 1980.

[12] Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, II. *Journal of Symbolic Computation*, 60:94–112, 2014.

[13] Christian Pech and Sven Reichard. *Enumerating Set Orbits*, pages 137–150. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.

[14] Leonard H. Soicher. On the structure and classification of somas: Generalizations of mutually orthogonal latin squares. *Electr. J. Comb.*, 6, 1999.